

**DARPA Urban Challenge
Technical Paper**



Submitted May 31st, 2007

Axion Racing
Team ID: B142
4607 Lakeview Canyon Road
Suite 428
Westlake Village, CA 91361
[Http://www.AxionRacing.com](http://www.AxionRacing.com)

Author:

William Kehaly
818-597-2703
BKehaly@AxionRacing.com

DISCLAIMER: The information contained in this paper does not represent the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency (DARPA) or the Department of Defense. DARPA does not guarantee the accuracy or reliability of the information in this paper.

Executive Summary

This technical paper details an introduction to Axion, LLC, an overview of problems to be solved, explains the design choices the team has made, and presents what will be the final realization of test and system performances. The final deliverable is a safe and effective autonomous vehicle able to perform in a dense urban area, utilize mission re-planning, road following, and obstacle avoidance, while interacting with other vehicles.

Axion, LLC has spent the last four years developing a superior autonomous vehicle design approach. Company has had complete success in both DARPA's 2004 QID and 2005 NQE autonomous vehicle trials. This success bodes well for a fully autonomous vehicle by November 2007.

Autonomous technology created by Axion, LLC, for the DARPA Urban Challenge, will reduce the cost and increase the availability of commercial-off-the-shelf (COTS) intelligent vehicle systems.

Scenarios for the fully functional autonomous vehicle includes going to a selected location, providing a safe round trip, parking itself after dropping of the vehicle's passengers, and will benefit both healthy and handicapped passengers.

It is Axion, LLC plan to produce an autonomous vehicle system that works exceedingly well in 2007 and will be the platform for autonomous developments in the future.

Table of Contents

Executive Summary	2
Table of Contents	3
1. Abbreviations	4
2. Axion Racing Introduction	5
2.1. Development Team Background.....	5
2.2. Support Team Background	7
2.3. UCSD Internships.....	7
2.4. Team Advisory Board	8
2.5. Previous Grand Challenge Performance.....	8
3. Introduction and Overview	9
3.1. System Architecture	9
4. Analysis and Design.....	10
4.1. Vehicle Selection.....	10
4.2. Power Source	10
4.3. Vehicle Control System.....	11
4.4. Computing and Network.....	11
4.5. Navigation.....	11
4.6. Sensors.....	11
4.7. Advanced Vision System.....	12
4.8. Reactive Sensory Enclosure.....	12
4.9. Laser Rangefinders.....	13
4.10. Axion Arbitrator.....	13
4.11. User Interface and Control System.....	14
5. Results and Performance	14
5.1. Sensor Integration and Interpretation	15
5.2. Decision-Making	17
5.2.1 Goal score creation	18
5.2.2 Goal Generation.....	18
5.3. Test Facilities	20
5.3.1. San Diego Streets.....	21
5.3.2. Torrey Pines Glider Port	21
5.3.3. Salton Sea, California	21
5.3.4. University of California, San Diego	22
5.3.5. Twenty Nine Palms WARTEC Facility.....	22
6. Conclusion	24

1. Abbreviations

The abbreviations used in this technical paper include the following:

Abbreviation	Definition
DARPA	Defense Advanced Research Projects Agency
DGC	DARPA Grand Challenge
ORCA	Online Representation and Certification Application
UGC	Urban Grand Challenge
RSE	Reactive Sensory Enclosure
IT	Information Technology

Table 1: Technical Paper Abbreviations

2. Axion Racing Introduction

Axion, LLC was established in 2003 to build autonomous vehicles. Company is ORCA registered and should be considered a non-traditional defense contractor for the Urban Challenge. The team will be racing in the 2007 DARPA Urban Challenge with no government funding and therefore is considered a “B” team. Axion Racing has proven able to qualify and race in both prior Grand Challenges.

Axion Racing’s Team Leader, and the team’s primary sponsor to this point, is based out of Westlake Village, California. The bulk of software and hardware modifications are conducted in San Diego, California. As a team, Axion Racing gets together weekly for status reports and updates of all aspects of autonomous vehicle developments. These meetings allow everyone to maintain a solid understanding of where the development process has taken the team.

Unlike college teams, Axion Racing is not affected by graduations and has lost only one member of the original team, by mutual agreement. This continuity has allowed for the team to concentrate on moving established success into continued beta and production applications.

The team has developed an excellent style of delivering working autonomous code without the help of government grants. With the core development team it is expected that Axion Racing will be able to not only meet, but exceed DARPA’s expectations for the 2007 Urban Challenge.

2.1. Development Team Background

The five members of Axion Racing development team are knowledgeable, personable, and work well together under the stress of creating working autonomous vehicles. It is expected that DARPA will be able to communicate its’ needs and requirements to the team, who will respond with deliverables that meet all requirements.

Name	Photo	Position	Responsibility
William Kehaly		Team Leader	All day-to-day activities during project. DARPA Point of Contact.
Clark Guest, PhD		Principal Investigator	Investigating that the team is always ahead of the autonomous curve.
Josh O’Briant		Chief Engineer	Blends base vehicle, software, and hardware, into urban vehicle.
Richard Kalling		Chief Information Officer	Works with team’s staff insuring timely code deliverables for vehicle.
Melanie Dumas		Software Engineer	Writes code and comments on code from other staff members.

Table 2: Axion Racing’s Development Team

William Kehaly (Team Leader)

Bill Kehaly has been the primary sponsor for Axion Racing's development of autonomous vehicle technology, over the last four years. He has a BA in Business Finance and a minor in Information Technology. Mr. Kehaly's first job was becoming the Turkey Accountant at Foster Poultry Farms. After two years, he took his computer skills and became an applications consultant with almost 100 different companies in the United States.

Along with IT consulting, Mr. Kehaly envisioned, managed development, and marketed a pen-based charting system for use by baseball and softball teams. This application was so successful, it was purchased by and he then consulted with seven Major League Baseball teams, USA Baseball, USA Softball, and nine division one colleges over nine years. He also spent a few years consulting with National Hockey League (Predators and Ducks) teams.

Mr. Kehaly's second full-time job was in the late 1990's, with eToys.com, where he successfully directed a \$1.5M European software implementation for the company. This was completed under budget and ahead of schedule, just prior to the company going bankrupt.

Dr. Clark Guest (Principal Investigator)

Dr. Guest is a professor in the Electrical and Computer Engineering department at the University of California, San Diego. He has directed research in optical computing, optical interconnections, pattern recognition, neural networks, and computer generated holography.

Dr. Guest's group was the first to provide a practical demonstration of the superiority of optical interconnections. They also have been the leading developers of simulated annealing used for the design of computer-generated holograms. Recently, Dr. Guest's group has been the first to successfully apply the Neocognitron vision model to grayscale images and binocular depth perception. His group was also responsible for developing the CAN neural network algorithm, which are more than four orders of magnitude faster than the widely-used back propagation.

Josh O'Briant (Chief Engineer)

Mr. O'Briant spends his time working on different types of robots. He has spent the past few years working in the electronics field. Many of these applications are developed for specific consumer usage by Josh's team of engineers. Raised in the Midwest, Josh spends much of his free time traveling and working on interesting applications, like the DARPA Grand Challenge.

As the Chief Engineer, Josh is responsible for programming and monitoring all of the specifics for the team's instrument installation and survivability. Josh has a team of support individuals that help him in his tasks.

Richard Kalling (Chief Information Officer)

Richard Kalling graduated from the University of Southern California with a BS in Computer Engineering and Computer Science. He is an original member of Axion Racing and is responsible for creating the team's path planning algorithms which takes into account terrain and road data and determines an optimal path.

Melanie Dumas (Software Engineer)

Ms. Dumas will participate with the software design and development, along overall system integration. Melanie was the lead programmer for the DGC 05 and DGC 06. She currently works for General Atomics Aeronautical Systems, developing software for the unmanned Predator aircraft.

2.2. Support Team Background

An additional seven members of the team have provided assistance and solid deliverables over the last four years. These individuals have provided this support, because that have both an equity share in Axion, LLC and they love the challenge of building safe autonomous vehicles.

Name	Photo	Position	Responsibility
Timothy Doyle		Vision Systems	Write code for vision camera systems and interact with Axion Arbitrator.
Greg Jones		EMC Engineer	Manages the electronic motion control installation and maintenance.
Jeremy Pollock		LADAR Systems	Chief investigator of LADAR applications and associated software development.
Walter Phillips		System Integration	Provides expertise on interaction of the different autonomous software and hardware.
Paul Kimmelman		Cartographer	Helps the team to plan multiple routes to reach a designated mission objective.
Joël Chenu		National Instruments	Creates software that allows the autonomous vehicle to utilize NI products to drive.
George Spalding, III		Components Engineering	Specializes in wiring and code for different sensors and systems in the autonomous vehicle.

Table 3: Axion Racing’s Support Team

2.3. UCSD Internships

As in prior DARPA Grand Challenges, it is expected that Axion Racing will have an abundance of University of California, San Diego students specializing in programming, robotics, and mechanical engineering interested in working on the team’s autonomous vehicle. Axion Racing has developed a good working relationship with the university and expects to continue to have lots of interest from students wanting to participate in autonomous vehicle design, through unpaid internships.

The internships provided by Axion Racing allow UCSD students to bring their class work and the team’s autonomous work together into a very valuable learning experience. This experience will provide the interns with the ability to be a part of this burgeoning industry.

2.4. Team Advisory Board

Over the years, the Axion Racing has been able to develop working relationships with many different types of individuals interested in helping the team overcome challenges to autonomous vehicle driving. These individuals make up the team's advisory board and are available to review many different types of issues and to provide excellent insight into answers for difficult questions that occasionally come up.

Name	Employer	Services	Available	Contact Information
Peter Williams, PhD	NavCom Technology	GPS Signals	As Needed	PWilliams@navcomtech.com
Mike Comberiate, PhD	NASA, Goddard Robotics	Autonomous Systems	Once a week	mcomberi@pop400.gsfc.nasa.gov
Ian Reynolds, PhD	Northrop Grumman, Inc	Inertial Navigation	As Needed	Ian.Reynolds@NGC.com
Lev Sadovnik, PhD	Waveband RADAR	MMW RADAR	As Needed	Lev.Sadovnik@sncorp.com
Vladimir Brajovic	Intrigue	Brightening road shadows	Once a week	Brajovic@Intriquetek.com
Renee Hendrick	SICK LADAR	Laser ranging system	As Needed	RenneH@Itsco.com
Martin Cox, PhD	BAE Systems, Inc	Autonomous Sensors	As Needed	MCoxforce@Yahoo.com
John F. Bower, Senior Scientist	US Marine Corp Pacific Experimentation Center	WARTEC Field Testing	As Needed	BowerJ@Battele.org
R. Douglas Ramsey, PhD	Utah State University	Mapping and routing	Once a week	Doug.Ramsey@usu.edu
James Stephens, PhD	University of Southern Mississippi	Physics and Inertia	As Needed	James.M.Stephens@usm.edu
Terrance Boulton, PhD	University of Colorado	Vision technology	As Needed	TBoulton@vast.uccs.edu
Serge Belongie, PhD	UC San Diego	Software Applications	Once a week	SJB@cs.ucsd.edu

Table 4: Team Advisory Board

2.5. Previous Grand Challenge Performance

Axion Racing's development team and support staff have been together for both the 2004 and 2005 DARPA Grand Challenges. This group has shown a consistent ability to meet the qualification needs detailed for both the QID and NQE held by DARPA, before each previous Grand Challenge. The team's experience will come in handy, while creating new algorithms and sensor fusion for autonomous vehicles to excel at the DARPA Urban Challenge. It is expected that the team's prior experience will provide an excellent platform to create useable autonomous vehicle systems. The team has also shown a distinct capability in creating technical papers that meet DARPA's requirements and provide knowledgeable

3. Introduction and Overview

Axion Racing proposes to develop a robotic vehicle for entry into the DARPA Urban Challenge that integrates the best of proven technology with innovative and unique ideas. The resulting robot is expected to perform at a high competitive level in the Challenge and provide a general purpose platform for future improvement and extension.

Axion Racing developed a very competitive vehicle, Spirit, for the first and second Grand Challenge events. Spirit took off as the sixth vehicle in the first event and was awarded fourth starting position, as a result of its performance in the 2005 NQE. The vehicle provided the seventh best performance in the 2005 Grand Challenge event, succumbing after 65 miles to a mechanical failure. All of this was accomplished by a team of engineers working after-hours from their full-time jobs, with a privately-funded budget that was an order of magnitude less than those of other successful teams. Axion Racing was the clear winner in terms of performance-per-dollar-spent.

For the DARPA Urban Challenge, Axion Racing proposes to build on their solid base of experience and unique technology already developed, while addressing the new issues presented by urban driving. Nearly all members of the previous team have committed to contribute to this effort. The team believes that experience has validated the LabView-based vehicle low-level control system, the multi-process–multi-machine arbitrated software architecture, and many of their sensor software and hardware interfaces.

Experience has also taught important lessons including the selection of low-power, low-heat processors and providing closed loop control of all actuator systems. The team has a solid plan for improving both the sensory and decision-making capabilities of their system to meet the challenges of driving in urban traffic.

3.1. System Architecture

The proposed Axion software system block diagram is shown in Figure 1. A convenient way of viewing it is in analogy to the human thought process: there is an unconscious object recognition and context mapping activity, and there is a conscious decision making activity based on the sensory information. There is also a path planning activity that is triggered by decisions and affects subsequent decisions. The final stage is closed-loop control system to supervise the implementation of decisions.

Axion, LLC has developed a highly competitive vehicle, which uses a superior system architecture, for the 2007 DARPA Urban Challenge.

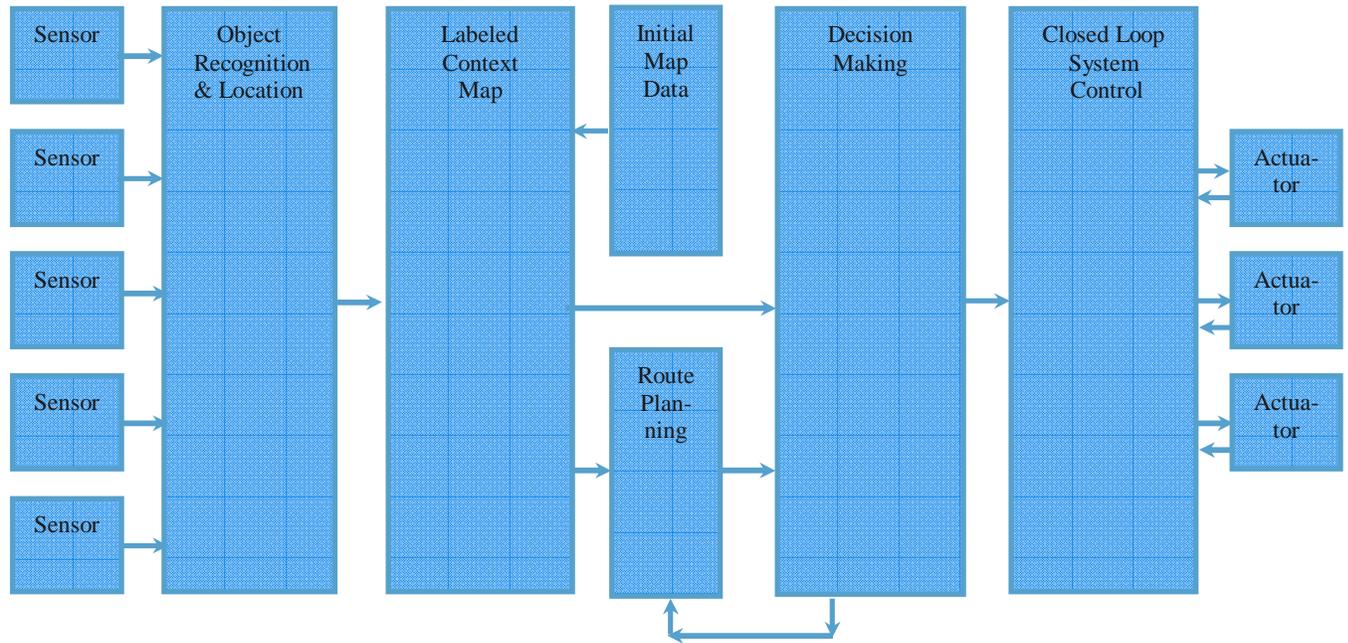


Figure 1: Software System Architecture

4. Analysis and Design

This section describes the answers to problems presented to Axion Racing by the 2007 DARPA Urban Challenge.

4.1. Vehicle Selection

Axion Racing selected a 1994 Jeep Grand Cherokee as its platform for the 2004 and 2005 Grand Challenge events. This vehicle is a reliable off-road performer, and it has the capability to transport people as well as a large payload. It satisfies both military and civilian requirements, including passenger comfort and hefty ruggedization. One of Axion's key objectives was to preserve the stock look and feel of the vehicle, in order to ease the public's acceptance of autonomous vehicles, and to convince potential vehicle manufacturers that the autonomous system will not significantly impact their existing designs.

4.2. Power Source

The autonomous vehicle power source must simultaneously power a multitude of sensors, several computers that analyze the sensor data, and still have enough amperage left to run the vehicle itself. The engineers of Axion Racing replaced the stock alternator with two higher power alternator units, delivering a total current of over 400 amps from a bank of four marine batteries. These batteries power two inverters, providing both AC and DC power for the system. The batteries alone are capable of powering the computer systems for several hours without recharging.

4.3. Vehicle Control System

The Axion Vehicle Control System consists of a National Instruments Compact Field Point real time control unit and a set of servos that manipulate the steering column, gas/brake unit, and the gear shifter. Commercial, off the shelf (COTS) products from Electronic Mobility Controls (EMC) operate the steering, accelerator, and brake. These products fit snugly beneath the steering wheel and against the floorboards, such that they do not impede manual operation. The real time controller sends analog signals to the motor controllers to position the steering and gas/brake servos. The real time controller consists of an embedded computer with a modular backplane for insertion of analog and digital I/O cards.



Figure 2: Vehicle Control System

The Vehicle Control System is also capable of driving the vehicle via remote control. Switching between manual and autonomous mode simply requires inserting or removing two pins.

4.4. Computing and Network

Five Dell 2650 dual Xeon Servers are shock mounted in a server rack in the luggage compartment of the vehicle. These servers have been operating for several years in rugged desert conditions, and are all covered under service contracts. Four servers run Windows 2000 to utilize API software for the cameras, and the primary control server runs Linux. A gigabit Ethernet connects the computers and the real time controller. Wireless networking provides external telemetry during test runs.

4.5. Navigation

The Axion platform is equipped with two NavCom Starfire GPS receivers. Data from these devices is relayed into a Northrop Grumman LN270 Inertial Navigation System. The LN270 is a cutting edge, precise navigation system that has not yet reached mass production. It uses a series of laser ring gyros to maintain accurate vehicle position in the event of a GPS outage. The integration of the Starfire network with the LN270 is the first of its kind, and produces navigational solutions more accurate than its market competitors.

4.6. Sensors

The Axion Sensor Suite consists of five stereo vision cameras, four Laser Detection and Ranging (LADAR) systems, and a high resolution color camera. Data from these sensors is independently evaluated, producing a “vote” that is sent to the Axion Arbitrator for processing. This section defines how the sensors are used, and the following section describes decision making process.

4.7. Advanced Vision System



Figure 3: Advanced Vision Systems

Stereo Vision

The showpiece of the Axion Advanced Vision System is a set of five Point Grey Research “Bumblebee” stereo cameras. These cameras produce a depth image of the area in front of the autonomous vehicle, with different lenses to achieve a near and far field of view. The stereo vision cameras measure the disparity between the two images in order to create a three dimensional image (this is similar to how human eyesight works).

The stereo images are processed by taking the depth map from the image, and breaking it into cells. The road is found by calculating the slope of each cell, and finding the cells with the smallest slope. Cells with large slopes are considered obstacles. Obstacle information is conveyed to the system by outputting a negative vote, which is stored in the Arbitrator.

Color Vision

The second component of the Advanced Vision System is a high resolution color camera from Allied Vision Technologies, the “Dolphin”. This camera produces 1600 x 1200 color images at 12 frames a second. The color camera frames are read by the Image Server, which performs a preprocessing step and passes the images to two color camera “voter” algorithms.

The first color camera voter algorithm uses texture to determine where the road lies. A road’s texture tends to have a vanishing point where the lines in the road converge, so this algorithm produces a vote to point the vehicle in that direction. The accuracy of finding roads using texture is over 95%. The second algorithm uses a neural network to find the road. Neural networks are very fast, since all the time to construct them occurs in a preprocessing step offline. The neural network is trained using a genetic algorithm, which finds optimal network weights from a set of training images.

4.8. Reactive Sensory Enclosure

The vision system must be protected from the elements, including rain, snow, mud, dust, and other particulates in an off-road environment that obscure each camera’s field of view. The Reactive Sensor Enclosure (RSE) protects the vision system through a sealed enclosure that can be mounted above the hood of many vehicles. On Axion’s vehicle, the assembly is incorporated into the design of the vehicle while out of the line of sight of a human driver. The cameras are protected by a sheet of tinted glass and a polarized film to reduce glare from the sun. The key feature of the RSE is a series of miniaturized wipers for each camera. A component of the

camera's processing algorithm is detection of particulate accumulation. If something is detected that obscures a camera's vision, the wipers are triggered for that particular camera.

4.9. Laser Rangefinders

Axion uses four laser ranging and detection systems (LADAR) from SICK Technologies. These systems boast impressive performance; 80 meter range, a 40Hz scan rate, and centimeter accuracy. Nearly every team in the DARPA Grand Challenge employed LADARs as their primary sensor.

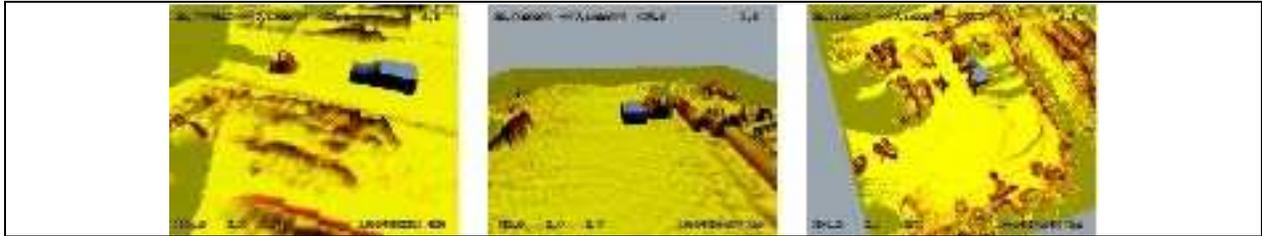


Figure 4: Local mapper output from LADAR data.

Axion Racing fused range data from its four LADAR sensors to create a single worldview. The key to this “local mapper” system correlates every laser scan is correlated with a point in world coordinates by combining the data with GPS coordinates. As the lasers scan in front of the vehicle, a very precise world map is built up. There is great potential for this map to be used to survey an area and provide high precision maps for other vehicles or logistics teams, with broad commercial appeal.

4.10. Axion Arbitrator

Each sensor provides positive, road-finding scores, or negative, obstacle-detection scores for all possible turning directions. The resulting scores are combined in the Axion Arbitrator, which uses its state machine and contextual information to choose the best path. The Arbitrator is a true plug and play system, which can adapt to sensors as they are added to and removed from the system. The Arbitrator evaluates the final scores, and decides speed and gearing as well as steering direction.

The Arbitration system is applicable to a wide range of ground vehicles, sensors, and computing platforms. The algorithm itself is very fast and can run on an embedded system or laptop. The system is generic; it accepts the standardized vote data from the sensors and applies a set of predetermined weights. The Axion Arbitrator technology is an ideal platform for prototyping sensor integration on unmanned vehicles.

One area of current research is the ability to detect contexts using the Arbitrator, which would adjust the weight of the voters as the vehicle moves through the different types of terrain. For example, the texture color camera algorithm performs well, and should have a high weight on a dirt road, but on a uniform, dry lakebed, the information it provides is negligible. In this case, a

context switch to reduce the influence this voter has on the system would improve road-finding performance. Detecting and adapting to changing contexts will be a key feature in the future of autonomous vehicles.

4.11. User Interface and Control System

The underlying complexity of Axion’s autonomous vehicle is simplified by a high level control user interface system. This system is capable of starting all applications quickly and displaying the results of changes immediately. On a touch screen mounted in the dashboard of the vehicle, the user can build a waypoint course, adjust the Arbitrator configuration, start and stop all components of the distributed system, and view the real time data as the autonomous vehicle chooses its course.

In the DARPA Grand Challenge, it was possible for vehicles to be stopped overnight and restarted by DARPA officials in the morning. As a result, starting the robot is a very simple procedure: on power up all the computers automatically launch their applications and await the “GO” signal from the official.

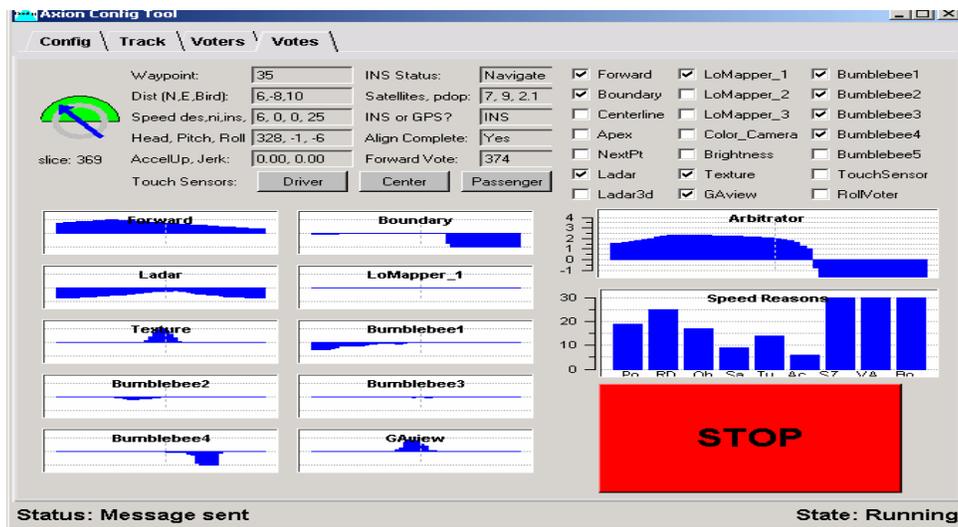


Figure 5: Primary user interface control screen

The system’s health is monitored with the user interface application, and the vehicle can be started simply by pressing the “GO” button.

5. Results and Performance

The following sections detail Axion Racing’s proposed Analysis and Design solution to the DARPA Urban Grand Challenge.

5.1. Sensor Integration and Interpretation

This section presents Axion's plans for sensor data acquisition and processing into information useful as input to driving decisions. First, candidate context sensor technologies will be discussed. Sensors fall into two broad categories: imaging and non-imaging sensors. A review is given of experience with sensors gained in previous Challenges, as well as new sensor types under consideration. An overview of combining data from non-imaging and imaging sensors is given. Then a detailed treatment of the process of translating multiple sources of imaging sensor data into a combined, consistent map of the vehicle context is presented. Finally, a plan for arranging context map information to serve as input to decision processing is described.

Sensors on an autonomous vehicle fall into three broad categories: position, status, and context. Position sensors include GPS and inertial navigation systems. These will provide input primarily to the route planning algorithm, but are largely unimportant to immediate driving decisions. Status sensors report on quantities internal to the vehicle such as steering angle, speed, or engine temperature. Some of these are useful in closed loop control of the vehicle systems and others for diagnostic purposes, but again they do not play a central role in decision making. One might argue, for example, that vehicle speed is important to driving decisions, but it should be set as an output of the decision process. If the set speed and the actual speed of the vehicle differ significantly, there is a failure in that control system. Context sensors report on the surroundings of the vehicle. They are the primary sources of information for navigational decisions. These sensors provide information about road boundaries, static obstacles, and moving vehicles around the robot.

Axion intends to use some context sensor types they have experience with from previous Challenges, as well as evaluate new sensor types made available by advancing technology, improved funding, or specific requirements of urban driving. Context sensor types used on Spirit in previous Challenges included scanning LADAR, stereoscopic video camera pairs, a color video camera, and some simple physical touch sensors. Experimentation was done with IR cameras and a RADAR sensor, but they were found to offer little incremental value for desert driving.

Axion intends to remain flexible to incorporate new sensor types during the development process until the hardware and software lockdown dates arrive. However, preliminary estimates of preferred sensor types have been prepared. Though quantitative performance data is not yet available, a new type of 360° scanning imaging LADAR available from Velodyne Systems is of great interest. Another new sensor type to evaluate is directional microphones to sense sound from nearby traffic. These two novel sensor types will be backed up by stereoscopic video camera pairs and color video cameras which both proved their worth in previous Challenges. IR cameras may assume increased importance if reliable IR signatures can be obtained from running vehicles. RADAR results were disappointing in previous challenges, but new types and vendors have appeared recently that merit evaluation.

Driving in urban traffic requires situational awareness of all four directions from the robotic vehicle. The 360° view of the scanning LADAR makes it particularly interesting for this reason. Supplementary long range sensors such as video cameras are appropriate for the forward and left and right forward quadrant directions; the latter directions for sensing crossing traffic when

preparing to merge. Mid-range and short-range sensors are appropriate for all vehicle directions to assess the proximity of other vehicles. The urge to deploy many sensor types in all directions must be balanced against not only the sensor cost, but also the processing power required for interpretation.

For context sensor types two broad categories are suggested: imaging and non-imaging. An example of a non-imaging sensor would be a directional microphone, while imaging sensors include LADAR and video cameras. Integrating data from both types in a meaningful way presents a particular challenge. We view non-imaging sensor data as both directive and corroborative. It is directive in the sense that it may be used to aid in the processing of labeling objects viewed by imaging sensors. For example, a running vehicle will have a different acoustic signature than a static vehicle. Non-imaging sensor data can be corroborative in the sense that it adds confidence to classifications produced by imaging sensors. For example, if an object is given an uncertain classification as a truck based on LADAR data, if the sound of a truck is coming from the same direction, the confidence in that classification is increased.

Many of the context sensor types will be imaging sensors, including LADAR, video cameras, stereoscopic video camera pairs, and even RADAR, if it is used. Some types of imaging sensors also provide direct indication of object range while others do not. The Axion plan is to fuse data from the various image sensors via an inverse mapping process. Initially, each imaging sensor will have an independent map of its surroundings. Images on the sensor detector will be geometrically projected back onto the map. For sensor types that do not provide range information, which positions ambiguity will be represented in the back projection. By prior calibration, the individual maps of all sensors will be aligned in registration. Data from the individual maps will be combined onto a master map. During this process, ambiguity can be resolved. Also, discrimination of significant and insignificant edges can be performed. For example, the edge of a shadow falling across a road as seen by a video camera is an insignificant edge, and can be removed by comparison with a LADAR image. The first step in developing an understanding of the robot vehicle context is the discrimination, noise filtering, and enhancement of edges obtained by fusing image data from multiple sensor types.

Another important source of information for image understanding is the history of images. Many edges in a scene can be expected to move in a predictable way due to movement of the robot through the scene. Objects that are in motion with respect to the ground will have their own predictable edge movements. Edges seen in previous images can be projected forward in time to new positions and help to further enhance current edge information.

The general case of connecting edge images through their motion over time is called optical flow computation. In some circumstances, general computation of optical flow can be computationally intensive, but numerous opportunities exist in the urban driving context to make it acceptably efficient. In addition to enabling enhancement of edges in a sequence of images, optical flow also offers information regarding the speed and direction of moving objects, their motion relative to the robot, their distance, and projecting whether they are on a collision course. We expect optical flow to be an important component of the image understanding process.

An intermediate goal of the image understanding process is to segment the image into objects such as road, obstacles, and other vehicles. Edges detected in earlier steps may not fully enclose each object because of noise, intervening closer objects, and so forth. To complete object boundaries, cues such as colinearity, color, motion, and texture can be used. Once objects are fully enclosed by edges, each one can be labeled. The labels will include not only the type of object, but also its motion, distance, and whether it's on a collision path.

Predicting the path of moving vehicles is a crucial task for success in urban driving. As mentioned above, computation of optical flow provides one measure of velocities. Comparing two sequential labeled context maps can provide another measure. Comparing three maps will yield information about acceleration which is important to assess braking and turning. Future positions of vehicles can be projected using simple physics formulae to predict whether they represent a potential collision hazard. A vehicle's future position may be represented not just as a single discrete location but as a probability distribution of positions the vehicle might achieve by changing course, with more extreme maneuvers assigned lower probability. Defensive driving strategy is to position the robotic vehicle where it has the lowest probability of intersecting any other vehicle.

One output of sensor data processing will be an actual two-dimensional map of the environment surrounding the robotic vehicle. This map will combine information known a priori such as road positions, speed limits, and stop sign locations with position sensor information to locate the robotic vehicle on the map, and context sensor information to show stationary obstacles and other moving vehicles. Maps can also show vehicle motions projected forward in time to identify potential collisions. Such graphical maps will be useful primarily as a resource for engineers developing the driving algorithms. For use by the robotic driving algorithms, other information formats will be more appropriate. Sensor location of roads can be used to cross check GPS and INS position information. Identified obstacles and vehicles can be presented as a list in priority order of likelihood and immediacy of potential collisions.

Most importantly, processed sensor information can be subjected to a series of tests to develop logical antecedent conditions for the decision-making processes that follow. For example, a simple antecedent would be "there is a vehicle traveling in the lane beside us". More complex antecedents can also be formed, such as "a vehicle is approaching on a crossing street and it is not braking". These antecedents are the culmination of the progression from raw sensor data to useful conceptual information.

5.2. Decision-Making

Axion will use a probabilistic approach to using the gathered world state information to make a decision on determining what routes are safe ones to take. Using a probabilistic approach has the advantage of allowing for both errors in sensor data and errors in predicting the future state of mobile objects. A route that is picked that has the highest probability of completing the given mission objectives in the least time while ensuring high levels of safety and obeying traffic laws. For the purposes of this paper, a route is defined as a series of points in time and space. After a route is determined, signals will be sent to the actuators so that actual following of the route can commence.

The first step in picking a route is testing to see if a route is a good one. In order to have the best possible planning potential paths will be continually tested. Only paths that are still relevant will be considered when choosing the path to follow. A route will be followed until a significantly better one has been discovered or changes to the world view render it invalid. It is good to keep in mind that routes can include moving through time as well as space. It is expected that at some point it will be unsafe to move in any direction, or moving backward would prove less effective than simply sitting still, in such a case the beginning of the route will move only through time. There are a very large number of paths that one can take so not all routes can be tested. As such, Axion will be using a heuristic to guide the path selection in the right direction. There are a wide variety of potential heuristics that can be applied to try and quickly find a good route. For instance, if the first path tried results in a very low score compared to a baseline average, a path that is very dissimilar could be tried next time. In addition, one could use a heuristic that first tested routes planned by a more simplistic algorithm using a subset of the world view, such as a route planned only on the information provided in the route network definition provided by DARPA, or if that data didn't exist, even a straight line path.

In order to accomplish this, Axion will be using a weighted set of goals. A goal is a measurable objective that directs the action of the vehicle. Dividing the decision making up into a set of goals makes the decision making modular. This has the advantage of making it easier to debug along with allowing for a larger number of people to concurrently work on the system. Modularity will be enforced by keeping the interface of the goals consistent. The interface will consist of gathering information that is provided by the world view and the score of other goals as input and output only their own score. Potential routes will have scores for each of the goals based on the probability of that route being able to fulfill that goal. The goals at the highest level will be weighted according to how valuable they are in maintaining safety and ensuring movement towards mission objectives, and will each contribute to the overall score of a particular route.

5.2.1 Goal score creation

Each potential path in time and space will be given a score corresponding to the likelihood that it will fulfill a given goal. These probabilities will be generated by information pulled the world view as well as the scores from other goals. For instance, if the goal is to maintain a safe following distance and the world view will provide the probability of the car being within the specified distance range at the time when the vehicle is in a projected position along a route, this will directly affect the score of this goal for this route. The other input for a goal is other goals; these are taken into account because the likelihood of success of one goal may be dependent on the scores of other goals. Due to processing time concerns, measures will be put in place to ensure that the depth of goal dependencies does not go down too far. The unresolved scores by not progressing further will be interpreted as uncertainty in the probability determinations.

5.2.2 Goal Generation

The set of goals that are used at the highest level will be based off of the rules that are outlined by DARPA for the Urban Challenge. For example, ensuring that there are no collisions or

stopping at a stop sign. In addition to the goals provided by DARPA, Axion will also include goals that encapsulate probability calculations, thus making it easier for a human monitoring the system during development to understand the reasoning behind the decisions that the vehicle makes. For example, a goal may be added to avoid steep inclines, as this has the effect of raising the probability that the vehicle is able to maintain speed and visibility, both of which may be other goals. In addition to developer created goals there will also be dynamically spawned goals. These goals will assist in the handling of high level planning as well as state dependent goals.

Some of the goals at the level directly contributing to the score of the path (as opposed to only contributing to the scores of other goals) will be goals that will have minimum likelihoods in order for a particular path to have any score at all. This allows for a higher level of insurance of safety and adherence to the Urban Challenge rules as well as being able to short circuit the rest of the subordinate rule calculations in order to save processing time. This includes goals like collision avoidance and feasibility analysis.

One of the goal generators will be the high level path planning module. The high level planner will use static elements in the world view in order to construct an optimal route to take to the next objective. Static elements include roads, terrain, and any objects that were discovered by the sensors that are unlikely to move. It will use a modified A* search to determine intermediate high level objectives. This allows the vehicle to intelligently determine which roads to follow or zones to traverse on the way to the next objective. For example, in the event that the vehicle needs to make a right turn after the next stop sign, it will generate a goal that has a high score if the route tested follows the waypoints in sequence that are specified for the current road in the direction of the stop sign, followed by the waypoints for the intersecting road to the right of the stop sign. While the actual stopping at the stop sign will be handled by other rules, it would be taken into account with regards to determining the optimal route.

In the case of the discovery of newly discovered road blockages, the world view will have been updated. The route planner will take note of this modified world view and modify the parts of the goal corresponding to the intermediate waypoints that are no longer valid. In addition, it will generate new waypoint goals that reflect this new world view. In this way previously unknown obstacles will not significantly affect the vehicles ability to proceed.

Goals may also be dynamically generated by state information. This is needed because there are many courses of action that are only useful if the action is completed. An example is when the vehicle comes to a dead end and can only proceed backward for a time before turning around. In order to prevent oscillation between forward and reverse it would be important that the vehicle go all the way back to a point where it is able to turn around before moving forward again. In this case a goal would be generated that would have a high score when the vehicle is going in a particular direction while in reverse. Another example of a state generated rule is when a 3 point turn is required. A rule would be generated that would get a high score when progress is made on rotating the vehicle by means of oscillating between forward and backward.

An important set of goals are the Feasibility goals. These goals get a high score when it is likely that the vehicle is actually capable of doing what was instructed. Examples of feasibility goals include being able to achieve the velocity required to be at a position at a particular time, if it is

possible to brake fast enough to avoid a collision with a vehicle instead of moving around it, if it is possible to go around a turn at the velocity indicated by the currently being tested path, along with a plethora of other real world checks.

Due to its high priority, special care is being given to collision avoidance. A set of goals will be created with the aim of not having collisions with other vehicles. A high score for these goals will occur when it is unlikely that the path chosen will result in one of the collision risk factors. The DARPA given rule for keeping a safe following distance is an example of one of the goals that would contribute to getting a high score in the collision avoidance rule (as well as contributing to the route score in its own right). The world view will be instrumental in providing the needed information for preventing collisions as it will provide probability regions for moving objects which the collision goal would use to calculate the likelihood of collision along the tested route.

Collision avoidance with regards to other vehicles may require more logic than simply using a physical model of the other vehicle. Objects with a high probability of being a car will be simulator to determine likely future positions. This simulator will be a simplified version of the AI that is used for its own navigation. This simulation will use its world view to map out probability regions for future states of a given vehicle. It would take into account known roads, stop signs and obstacles. This would then be used to check for probable intersections, and the score for any route would be lowered that goes over these regions would have their collision probabilities raised, and thus collision avoidance scores lowered accordingly.

After a route has been planned that is both safe and feasible, the needed signals are sent to the actuators to set the vehicle along that course. The actual amounts of steering adjustment, gas, and break applied will also be probabilistic in nature, applied at appropriate amounts to get the highest probability of achieving desired velocities and locations at appropriate moments in time. These values will also be determined by the world view and include things like the current incline, traction and the effectiveness of previous application of throttle.

The modular nature of the design will prove valuable during tweaking and debugging phase of development in that it lends itself way to the display of data. Axion will create a visualization application that will display the world view and annotate it with the scores of selected goals over explored paths. In addition to displaying explored paths, a developer will be able to suggest a path to explore and the system will calculate the goal scores and graphically display the results. Calculations based on input from other rules will be able to be adjusted on the fly. Another view of decision information will be an overall listing of the scores that resulted in the selection of the final chosen route, in order of their contribution to that chosen route. These tools will allow Axion to quickly iterate the vehicle and provide good indications as to what vehicle tweaks should be implemented next.

5.3. Test Facilities

Axion Racing tests autonomous code in the streets of San Diego, the Torrey Pines Glider Port, Salton Sea, and the University of California, San Diego campus. These test facilities provide the team with areas that mimic specific areas of interest for the team. The team has also received preliminary approval that it will be allowed to use the Twenty Nine Palms WARTEC facility

after a successful appointment to participate in the 2007 DARPA Urban Challenge. Axion Racing was one of the few California based teams to utilize this excellent facility in 2005.

5.3.1. San Diego Streets

Axion Racing takes input from production and test sensors, while the team drives Spirit around city streets of San Diego. All input is used to check the validity of each sensor, working in conjunction with other sensors and the Axion Arbitrator.



Figure 6: San Diego Streets

5.3.2. Torrey Pines Glider Port

This facility was used in 2005 for the DARPA Site Visit and continues to be used by Axion Racing to test code and vehicle adjustments. The team has been able to establish an excellent working relationship with the owners of the glider port and expect to continue testing there as long as needed.



Figure 7: Torrey Pines Test Facility

This facility provides a safe and easily accessible spot for all of the team to get together and test autonomous vehicle code. Located close are all services required to meet the needs of the team.

5.3.3. Salton Sea, California

Axion Racing will spend a few weekends in the Salton Sea area of the Mohave Desert. The team has spent extensive testing in this and the Barstow area and enjoys the opportunity to stress test

the vehicle algorithms under extreme conditions. Using this test area provides the team with open areas and excellent conditions.



Figure 8: Salton Sea Test Area

This location is very convenient for Axion Racing’s San Diego based team and provides many different types of environments to use.

5.3.4. University of California, San Diego

The team has been in contact with the manager of the student parking lot and he has agreed to let Axion Racing use the Gilman parking facility to test parking algorithms. Using this facility provides the ultimate setup for making sure that the team’s autonomous vehicle can park without incidents.

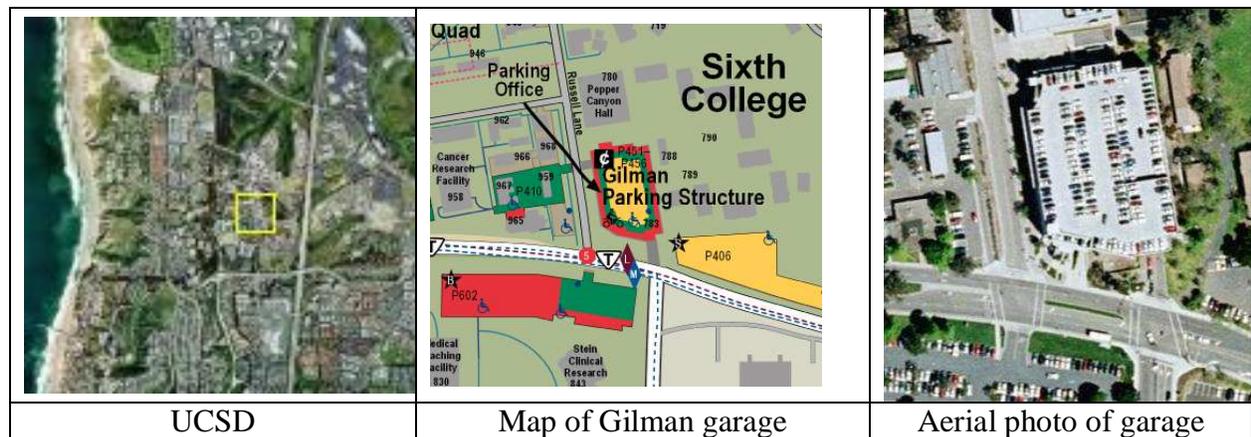


Figure 9: University of California, San Diego

It is planned that the team will be provided two floors to test during weekends and it will be Axion Racing’s responsibility to insure that non-team vehicles are prevented from entering any test area.

5.3.5. Twenty Nine Palms WARTEC Facility

The Twenty Nine Palms WARTEC Marine Base facility was visited once by Axion Racing during preparations for the 2005 DGC. The team has been in communication with the support

staff and it has been verified that they feel confident that we can use this facility in preparing for the 2007 Urban Challenge.



Figure 10: Twenty Nine Palms Test Facility

The WARTEC facility is especially nice because it sits in a large US Marine base and provides most of the scenarios required for testing a useful autonomous vehicle. Axion Racing finds this particular facility to be one of the most inviting government areas in the four years of autonomous racing experience.

6. Conclusion

Axion, LLC believes that it has the experience, expertise, and plan to win the 2007 DARPA Urban Challenge. By following the plan, detailed in this technical paper, the team expects to meet and exceed its' successes in the 2004 and 2005 DARPA Grand Challenge.

It is Axion Racing's plan to compete with and surpass the capabilities of all eleven teams that were provided one million dollars of government funding, under track "A" solicitation number BAA06-36, for the 2007 DARPA Urban Challenge.

Thank you for taking the time to review this technical paper.